arXiv:1405.0789v3 [cs.DM] 19 Aug 2014

A note on the ring loading problem*

Martin Skutella†

Abstract

The Ring Loading Problem is an optimal routing problem arising in the planning of optical communication networks which use bidirectional SONET rings. In mathematical terms, it is an unsplittable multicommodity flow problem on undirected ring networks. We prove that any split routing solution to the Ring Loading Problem can be turned into an unsplittable solution while increasing the load on any edge of the ring by no more than $+\frac{19}{14}D$, where D is the maximum demand value. This improves upon a classical result of Schrijver, Seymour, and Winkler (1998) who obtained a slightly larger bound of $+\frac{3}{2}D$. We also present an improved lower bound of $\frac{11}{10}D$ (previously $\frac{101}{100}D$) on the best possible bound and disprove a famous long-standing conjecture of Schrijver et al. in this context.

1 Introduction

An instance of the Ring Loading Problem is given by an undirected ring (cycle) on node set $V = \{1, 2, \ldots, n\}$ (numbered clockwise along the ring) with demands $d_{i,j} \geq 0$ for each pair of nodes $i < j$ in V. The task is to route all demands unsplittably, that is, each demand $d_{i,j}$ needs to be routed from node i to node j either in clockwise direction on the path $i, i+1, \ldots, j$ or in counterclockwise direction on the path $i, i-1, \ldots, j$. The objective is to minimize the maximum load on an edge of the ring. The optimum solution value is denoted by L. A small example instance of the Ring Loading Problem is given in Figure 1.

Application in telecommunication networks. This optimization problem arose in the early 1990ies as a crucial subproblem in the design of survivable telecommunication networks utilizing fiber-optic-based technologies; it was first studied by Cosares and Saniee [1] who also introduced the name Ring Loading Problem.

Fiber-optic-based telecommunication networks are built according to a standard called SONET (Synchronous Optical Networks) or the later derived SDH (Synchronous Digital Hierarchy) standard. These networks are based on rings of nodes (switches) with optical links between each pair of neighboring nodes. The actual bandwidth available along any edge of a ring is essentially determined by the capacity of the add/drop multiplexers (ADMs) installed at its nodes which make

*Large parts of the research described in this paper were conducted while the author was visiting the University of British Columbia (UBC) in Vancouver, Canada. It was partially supported by the DFG Research Center Matheon "Mathematics for key technologies" in Berlin, by DFG Priority Programme 1307 "Algorithm Engineering" (grant SK 58/6-3), by the Alexander von Humboldt-Foundation, and by the Peter Wall Institute for Advanced Studies at UBC.

†TU Berlin, Institute of Mathematics, MA 5-2, Straße des 17. Juni 136, 10623 Berlin, Germany, martin.skutella@tu-berlin.de

Figure 1: Left: An instance of the Ring Loading Problem on a ring of $n = 4$ nodes and two nonzero demands $d_{1,3} = d_{2,4} = 2$. Right: An unsplittable routing where the first demand is routed counterclockwise from node 1 to node 3, and the second demand is routed clockwise. The maximum edge load is 4 and this is obviously optimal, i. e., $L = 4$.

up a large part of the hardware cost. An important goal in network planning is thus to minimize the bandwidth required to satisfy all demands. In the early 1990ies Bell Communications Research (Bellcore) developed a network planning software called the 'SONET Toolkit' which also initiated the scientific study of the Ring Loading Problem. For a more detailed account of the technical background and history, we refer to the articles [1, 13, 8, 14].

Known results. Cosares and Saniee [1] prove by a reduction from the problem Partition that the Ring Loading Problem is weakly NP-hard. It is not known whether the problem is even strongly NP-hard or can be solved in pseudo-polynomial time. If all non-zero demands are equal, the problem can be solved in polynomial time; this follows from the work of Frank [4] (see also [5]) and is based on a theorem of Okamura and Seymour [11].

The hardness of the general problem motivates the consideration of the relaxed version of the Ring Loading Problem where demands may be split, i. e., a demand can be sent partly clockwise, partly counterclockwise. The optimum ring load L^* of a split routing is obviously a lower bound on the optimum load L of an unsplittable routing. Cosares and Saniee [1] prove that $L \leq 2L^*$, and this inequality is tight; see, e. g., the instance depicted in Figure 1 where $L = 4$ and $L^* = 2$. Moreover, they present an efficient algorithm for obtaining an unsplittable routing of load at most $2L^*$. For the problem of finding an optimum split routing, Myung, Kim, and Tcha [10] give a clever combinatorial algorithm with running time $O(nk)$, where k is the number of nonzero demands.

Schrijver, Seymour, and Winkler, in a landmark paper [12] (see also the annotated reprint [13]), present a clever analysis for a simple greedy algorithm that turns any split routing into an unsplittable routing while increasing the load on any edge by at most $\frac{3}{2}D$, where $D := \max_{i,j} d_{i,j}$ is the maximum demand value. Their result thus implies $L \leq L^* + \frac{3}{2}D$. Khanna [8] shows how to use this bound in order to obtain a polynomial-time approximation scheme for the Ring Loading Problem. For the special case of uniform nonzero demands $d_{i,j} = D$, Schrijver et al. even prove $L \leq L^* + D$.

On the other hand, Schrijver et al. exhibit an instance of the Ring Loading Problem together with a carefully chosen split routing that cannot be turned into an unsplittable routing without increasing the load on some edge by at least $\frac{101}{100}D$. This observation, however, does *not* imply a gap strictly larger than D between the optimum values of split and unsplittable routings. In the conclusion of their paper, Schrijver, Seymour, and Winkler write:

> "... even though the mathematics refuses to cooperate, *we* guarantee $L \leq L^* + D$."

In his excellent survey [14], Shepherd restates this 'guarantee' as Conjecture 1 and indicates its 'tempting similarity' to a result of Dinitz, Garg, and Goemans [3] on single-source unsplittable flows and a related conjecture of Goemans [2] (see also [15, 9]).

Our contribution. Our main result is the following theorem.

Theorem 1. *Any split routing solution to the* Ring Loading Problem *can be turned into an unsplittable routing while increasing the load on any edge by no more than* $\frac{19}{14}D$.

In particular, this result implies $L \leq L^* + \frac{19}{14}D$. Our algorithm runs in linear time and combines pairs of solutions obtained by the greedy algorithm of Schrijver, Seymour, and Winkler [12] using a clean crossover operation. In the area of genetic algorithms, it is known that the additional use of crossover (also called 'recombination' or 'sexual reproduction') is empirically more effective than mutation only, and this can sometimes even be confirmed by theoretical evidence; see, e. g., [7, 6]. We are, however, not aware of previous approximation results relying on such crossover operations.

We also exhibit a relatively simple instance of the Ring Loading Problem together with particular split routings that cannot be turned into unsplittable routings without increasing the load on some edge by at least $\frac{11}{10}D$. After more than 15 years, our results are the first improvements on the classical results of Schrijver, Seymour, and Winkler [12] (upper bound $\frac{3}{2}D$ and lower bound $\frac{101}{100}D$). Last but not least, we present an instance with $L = L^* + \frac{11}{10}D$, thus disproving Schrijver et al.'s long-standing conjecture $L \leq L^* + D$.

We finally mention an interesting combinatorial implication of our results. Schrijver et al. [12] define β to be the infimum of all $\alpha \in \mathbb{R}$ such that the following combinatorial statement holds: for all positive integers m and nonnegative reals $u_1, \ldots, u_m$ and $v_1, \ldots, v_m$ with $u_i + v_i \leq 1$, there exist $z_1, \ldots, z_m$ such that for every k, $z_k \in \{v_k, -u_k\}$ and

$$\left| \sum_{i=1}^{k} z_i - \sum_{i=k+1}^{m} z_i \right| \leq \alpha \ .$$

Schrijver et al. prove that $\beta \in [\frac{101}{100}, \frac{3}{2}]$. As a consequence of our results, we narrow this interval down to $\beta \in [\frac{11}{10}, \frac{19}{14}]$.

Outline. In Section 2 we present preliminary observations and results. In Section 3 we show how to improve the bound of $\frac{3}{2}D$ in case there happens to be a demand of 'medium size'. Subsequently, in Section 4, we deal with the remaining instances and complete the proof of Theorem 1. Finally, in Section 5 we present the improved lower bound as well as the counterexample to Schrijver et al.'s conjecture.

2 Preliminaries

In this section we present preliminary observations which can already be found – explicitly or implicitly – in the work of Schrijver et al. [12]. Two demands $d_{i,j}$ and $d_{k,\ell}$ are said to be *parallel* if there is an i-j-path and a k-ℓ-path that are edge-disjoint; otherwise they are said to *cross*.

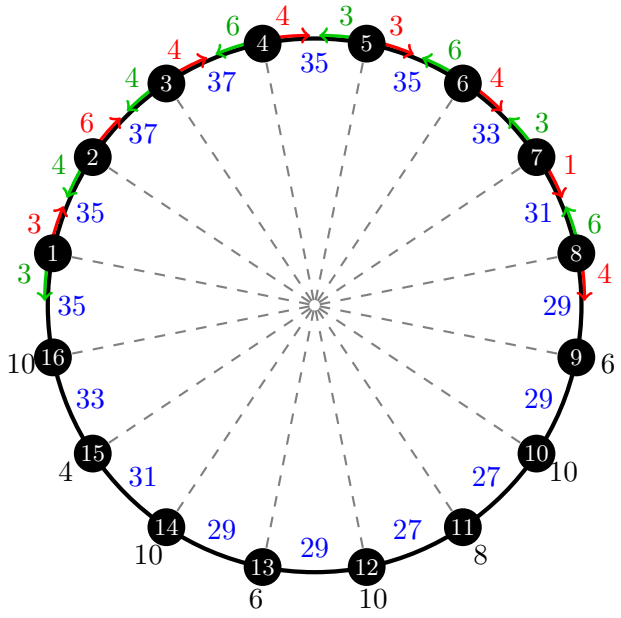

Figure 2: An instance of the Ring Loading Problem on $n = 16$ nodes with $m = 8$ nonzero demands for opposite node pairs and a split routing. There is, for example, a demand of value $d_7 = d_{7,15} = 4$ of which $u_7 = 1$ unit is routed clockwise and $v_7 = 3$ units are routed counterclockwise from node 7 to node 15. The maximum demand value is $D = 10$ and the maximum edge load is 37 (on edges $\{2, 3\}$ and $\{3, 4\}$).

Observation 1. *For an arbitrary split routing of two parallel demands, flow can be rerouted such that one of the two demands is routed unsplittably without increasing the load on any edge.*

Proof. For $q = 1, 2$, let P_q^a, P_q^b be the two paths of the q^{th} demand, and let x_q^a and x_q^b be the flow routed along these paths, respectively. Since the two demands are parallel, we may assume that P_1^a and P_2^a are edge-disjoint. Notice that P_1^b thus contains P_2^a, and P_2^b contains P_1^a. Increasing flow on P_1^a and P_2^a by $\min\{x_1^b, x_2^b\}$ and decreasing flow on P_1^b and P_2^b by the same amount therefore yields the desired routing. □

In order to prove Theorem 1, demands that are unsplittably routed in the given solution can be ignored (deleted). Thus, as a consequence of Observation 1, it suffices to prove the theorem for instances without parallel demands. In particular, any node is endpoint of at most one demand. Moreover, if some node of the ring is not an endpoint of any demand, its two incident edges always have the same load such that they can be merged and their common node be deleted.

After this preprocessing, we are left with a ring of even length $n = 2m$ and m demands $d_i := d_{i,i+m} > 0$, $i = 1, 2, \ldots, m$. Moreover, a split routing is given where, for each $i = 1, \ldots, m$, exactly $u_i > 0$ units of flow are sent from i to $i + m$ in clockwise direction and the remaining $v_i := d_i - u_i > 0$ units are sent counterclockwise; an example is given in Figure 2.

An unsplittable solution to the Ring Loading Problem can be represented by an m-tuple $z = (z_1, z_2, \ldots, z_m)$ with $z_i \in \{v_i, -u_i\}$ for $i = 1, \ldots, m$. Here, $z_i = v_i$ means that demand i is routed clockwise, i. e., compared to the given split routing, the flow on this path is increased by v_i while the flow on the counterclockwise path is decreased by v_i. Correspondingly, $z_i = -u_i$ means

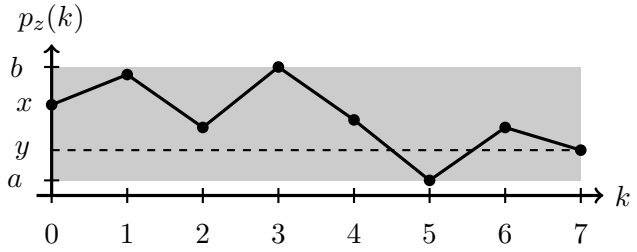

Figure 3: Pattern p_z living on strip $[a, b]$ with start point x, end point y.

that demand i is routed counterclockwise such that the flow on the clockwise path is decreased by u_i while the flow on the counterclockwise path is increased by u_i. As a consequence, for $k = 1, \ldots, m$, the load on edge $\{k, k+1\}$ changes by

$$\sum_{i=1}^{k} z_i - \sum_{i=k+1}^{m} z_i \ ,$$

and the load on the opposite edge $\{k+m, 1+(k+m \mod 2m)\}$ changes by the negative of this amount. Thus, compared to the given split routing, the maximum increase of flow on any edge is

$$\max_{k=1,\ldots,m} \left| \sum_{i=1}^{k} z_i - \sum_{i=k+1}^{m} z_i \right| \ .$$

We refer to this value as the *additive performance* of solution z.

For the sake of a more intuitive understanding, we use the following alternative representation of an unsplittable solution z: For an arbitrary but fixed $x \in \mathbb{R}$, let

$$p_z(k) := x + \sum_{i=1}^{k} z_i \qquad \text{for } k = 0, 1, \ldots, m.$$

An intuitive graphical representation of p_z is the graph of the continuous function on $[0, m]$ obtained from p_z by linear interpolation; see Figure 3 for an illustrating example. Notice that

$$p_z(k) - p_z(k-1) = z_k \in \{v_k, -u_k\} \qquad \text{for } k = 1, \ldots, m.$$

We refer to p_z as a *pattern starting at* $x = p_z(0)$ *and ending at* $y = p_z(m)$. Let $a := \min_{k=0,\ldots,m} p_z(k)$ and $b := \max_{k=0,\ldots,m} p_z(k)$. In reference to Figure 3 we refer to the interval $[a, b]$ as a *strip* and say that the *pattern lives on strip* $[a, b]$ *of width* $b - a$.

In the remainder of this paper we use the terms 'pattern' and 'unsplittable solution' interchangeably. Notice that an unsplittable solution corresponds to several patterns as the start point x can be arbitrarily chosen. All these patterns are, however, identical up to vertical shifting. To make the correspondence between unsplittable solutions z and patterns p_z unique, one could, for example, require $x = 0$ or $a = 0$. As these requirements would cause unnecessary technicalities in the presentation, we rather live with the ambiguity which doesn't cause any further problems.

Observation 2. *Let z be an unsplittable solution and consider a corresponding pattern p_z living on strip $[a,b]$ with start point x and end point y.*

(i) The unsplittable solution z and its pattern p_z have additive performance

$$\max_{k=1,\dots,m} \left| \sum_{i=1}^{k} z_i - \sum_{i=k+1}^{m} z_i \right| = \max\{2b - x - y, x + y - 2a\} \ . \tag{1}$$

(ii) The additive performance of pattern p_z is at least the width $b-a$ of strip $[a,b]$. This performance is achieved if and only if $x+y = a+b$, i. e., start point x and end point y are symmetric with respect to the center of interval $[a,b]$.

(iii) The additive performance of pattern p_z is at most $b-a+\varepsilon$ if and only if the pattern starts at x and ends at $y \in [y_{opt}-\varepsilon, y_{opt}+\varepsilon] \cap [a,b]$ with $y_{opt} := a+b-x$.

Proof. Part (i) follows from the fact that the maximum on the left hand side of (1) is obtained at an index k for which $p_z(k)$ is either maximal or minimal. The remaining statements follow immediately. □

For any $a' \in \mathbb{R}$ and $x \in [a', a'+D]$, one can easily construct a pattern starting at x and living on a strip $[a,b] \subseteq [a', a'+D]$ of width at most D by iteratively applying the following trivial observation.

Observation 3. *Let $d = u + v$ with $u, v \geq 0$. If I is an interval of size at least d and $x \in I$, then $x + v \in I$ or $x - u \in I$ (or both).*

Choosing for example $a' = 0$ and the start point $x = \frac{1}{2}D$, Schrijver et al. [12] construct an unsplittable solution with additive performance at most $\frac{3}{2}D$. In our terminology, their performance bound follows from Observation 2 (iii).

We notice that the described approach for constructing a pattern living on a strip of width at most D can also be applied in backwards direction. In this way, for any $a' \in \mathbb{R}$ and $y \in [a', a'+D]$, one gets a pattern ending at y and living on a strip $[a,b] \subseteq [a', a'+D]$.

3 Dealing with demands of medium size

It is not difficult to observe that the additive performance of $\frac{3}{2}D$ can be slightly improved in case there is a demand value of 'medium size'.

Lemma 1. *Let $\delta \geq 0$. If there is a demand i of value $d_i \in [\delta D, (1-\delta)D]$, then there exists a pattern with additive performance at most $\left(\frac{3}{2} - \frac{\delta}{2}\right)D$.*

Proof. By reindexing the nodes of the ring accordingly, we may assume $i = m$ without loss of generality. Deleting demand m (e. g., setting its demand value to $d_m := 0$) yields a smaller instance. For this smaller instance we consider a pattern living on a sub-strip of $[0,D]$, ending at $\frac{1}{2}(D+d_m) - v_m \in [0,D]$, and starting at some $x \in [0,D]$. This pattern can be extended to a pattern for the original instance including demand m in two different ways: the first (choosing $z_m = v_m$) ending at $\frac{1}{2}(D+d_m)$ and the second (choosing $z_m = -u_m$) ending at $\frac{1}{2}(D-d_m)$. By construction, both patterns start at the same point $x \in [0,D]$ and live on a sub-strip of $[0,D]$; see Figure 4 for an illustration.

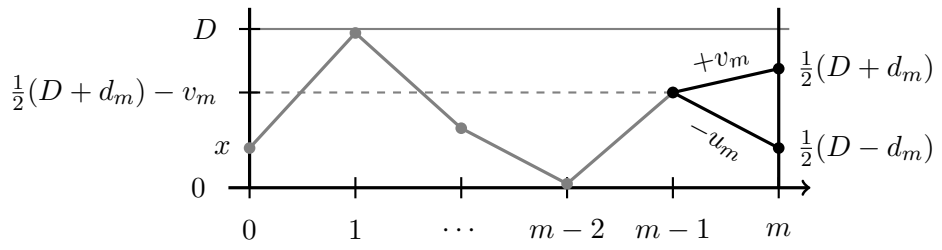

Figure 4: Illustration of the two patterns constructed in the proof of Lemma 1.

It follows from Observation 2 (i) that the additive performance of the first pattern ending at $\frac{1}{2}(D+d_m)$ is at most

$$\max\{2D - x - \tfrac{1}{2}(D+d_m), x + \tfrac{1}{2}(D+d_m)\} \ .$$

Since $x \geq 0$ and $d_m \geq \delta D$, the first of the two terms is at most $\left(\frac{3}{2} - \frac{\delta}{2}\right)D$. Moreover, if $x \leq \frac{1}{2}D$, then the same bound holds for the second term as $d_m \leq (1-\delta)D$. Finally, if $x \geq \frac{1}{2}D$, the additive performance of the second pattern is at most $\left(\frac{3}{2} - \frac{\delta}{2}\right)D$ by symmetric arguments. □

4 Instances with small and big demands only

In this section we show how to obtain an improved additive performance for instances without medium size demand values. More precisely, for some $0 < \delta < \frac{1}{2}$ to be determined later, we consider instances such that

$$d_i \in [0, \delta D] \cup [(1-\delta)D, D] \qquad \text{for every demand } i.$$

The next definition will turn out to be useful in order to identify pairs of patterns that can easily be combined to a new pattern while keeping the width of the strip it is living on reasonably small.

Definition 1. *Let $\varepsilon \geq 0$. Two patterns $p_{z'}$ and $p_{z''}$ are said to be* ε-close *if $|p_{z'}(k) - p_{z''}(k)| \leq \varepsilon$ for some $k \in \{0, 1, \ldots, m\}$.*

The following lemma defines a useful crossover operation for two ε-close patterns.

Lemma 2. *Consider a fixed instance. Let $p_{z'}$ be a pattern with start point x' living on strip $[a', b']$, and $p_{z''}$ a pattern with end point y'' living on strip $[a'', b'']$. If the two patterns are ε-close for some $\varepsilon > 0$, then there is a pattern p_z living on a sub-strip of*

$$\left[\min\{a', a''\} - \tfrac{1}{2}\varepsilon, \max\{b', b''\} + \tfrac{1}{2}\varepsilon\right] \tag{2}$$

with start point x and end point y such that $x + y = x' + y''$.

Proof. Let k be chosen according to Definition 1 and let $\varepsilon' := p_{z'}(k) - p_{z''}(k)$ such that $|\varepsilon'| \leq \varepsilon$. Define solution z by $z_i := z'_i$ for $i = 1, \ldots, k$ and $z_i := z''_i$ for $i = k+1, \ldots, m$. The corresponding pattern p_z with start point $x := x' - \frac{1}{2}\varepsilon'$ has end point $y = y'' + \frac{1}{2}\varepsilon'$ and lives on a sub-strip of (2); see Figure 5 for an illustration. □

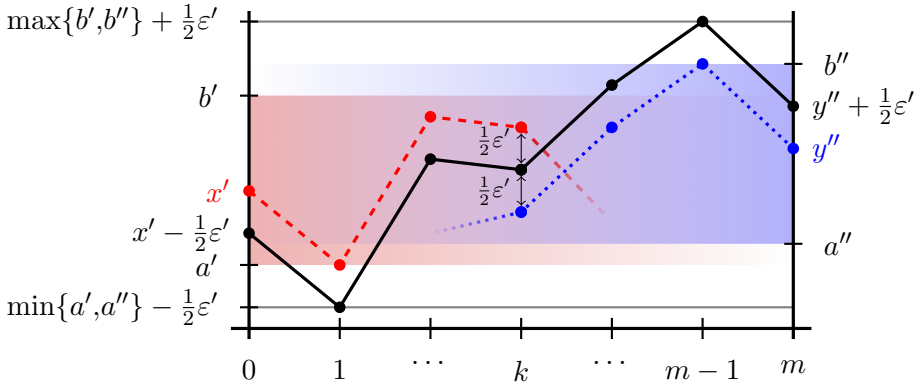

Figure 5: Illustration of the crossover operation used in the proof of Lemma 2.

As discussed in Section 2, it follows from Observation 3 that for a given start point $x = p_z(0) \in [0, D]$ a pattern z living on a sub-strip of $[0, D]$ can be constructed by iteratively choosing $z_k \in \{v_k, -u_k\}$ such that

$$p_z(k) = p_z(k-1) + z_k \in [0, D], \qquad \text{for } k = 1, \ldots, m.$$

Notice that the choice of z_k in this procedure is not always unique since both $p_z(k-1) + v_k$ and $p_z(k-1) - u_k$ may lie in strip $[0, D]$. In these situations a natural choice is to set z_k such that $p_z(k)$ is as close as possible to the midpoint $D/2$ of strip $[0, D]$, that is, as far as possible away from the boundaries 0 and D. A pattern constructed according to this refined procedure (breaking remaining ties arbitrarily) is called a *forward greedy pattern*. For technical reasons we also assume that the start point of a forward greedy pattern is far enough away from the borders of the strip $[0, D]$, that is $p_z(0) \in \left[\frac{1}{4}\delta D, (1 - \frac{1}{4}\delta)D\right]$.

Similarly, for a given end point $y = p_z(m) \in \left[\frac{1}{4}\delta D, (1-\frac{1}{4}\delta)D\right]$, a *backward greedy pattern* z living on a sub-strip of $[0, D]$ is constructed by iteratively choosing $z_k \in \{v_k, -u_k\}$ and setting $p_z(k-1) = p_z(k) - z_k$ such that $|p_z(k-1) - \frac{1}{2}D|$ is minimal, for $k = m, \ldots, 1$. We say that z is a *greedy pattern* if it is a forward or backward greedy pattern according to the definitions above.

Lemma 3. *Consider two greedy patterns p_z and $p_{z'}$. If there is a $k \in \{1, \ldots, m\}$ such that*

(i) $d_k \geq (1-\delta)D$ and $z_k \neq z'_k$,

(ii) $p_z(k-1) \geq p_{z'}(k-1)$ and $p_z(k) \geq p_{z'}(k)$,

then p_z and $p_{z'}$ are $\frac{1}{2}\delta D$-close.

An illustration of a situation described in the lemma is given in Figure 6 (a). Even without the assumption that p_z and $p_{z'}$ are greedy patterns, it is not difficult to derive from Figure 6 (a) a weaker version of this lemma that only guarantees δD-closeness (remember that $d_k = v_k + u_k \leq (1-\delta)D$). In general, however, the stronger bound holds for the case of greedy patterns only.

Proof. We consider the scenario where $z_k = v_k$ and thus $z'_k = -u_k$ (see Figure 6 (a)). The other scenario (i. e., $z_k = -u_k$ and $z'_k = v_k$) can be transformed into our scenario by reading the patterns p_z and $p_{z'}$ backwards from m to 0 (right to left). In particular, arguments symmetric to those given below can be used.

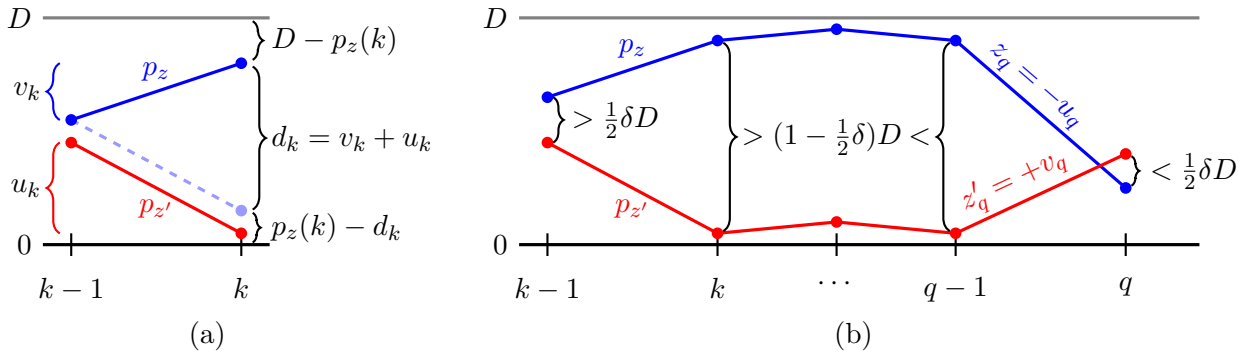

Figure 6: (a) Two patterns p_z and $p_{z'}$ featuring the properties stated in Lemma 3. Here we assume that $z_k = v_k$ and $z'_k = -u_k$. The other possibility ($z_k = -u_k$ and $z'_k = v_k$) can be illustrated by mirroring both patterns left to right. (b) Illustration of *Case III* in the proof of Lemma 3.

Case I: If p_z is a *forward* greedy pattern, then $p_z(k) - d_k \leq D - p_z(k)$ (otherwise, during the construction of the forward greedy pattern p_z, one would have chosen $z_k = -u_k$; see Figure 6 (a)). This yields

$$\begin{aligned} 0 \leq p_z(k-1) - p_{z'}(k-1) &= p_z(k) - v_k - \big(p_{z'}(k) + u_k\big) \leq p_z(k) - (v_k + u_k) \\ &= \tfrac{1}{2}\big(p_z(k) - d_k\big) + \tfrac{1}{2}\big(p_z(k) - d_k\big) \leq \tfrac{1}{2}\big(D - p_z(k)\big) + \tfrac{1}{2}\big(p_z(k) - d_k\big) = \tfrac{1}{2}(D - d_k) \leq \tfrac{1}{2}\delta D \ . \end{aligned}$$

Case II: Analogously, if $p_{z'}$ is a *forward* greedy pattern, then $D - p_{z'}(k) - d_k \leq p_{z'}(k)$ (otherwise, during the construction of the forward greedy pattern $p_{z'}$, one would have chosen $z'_k = v_k$). This yields

$$\begin{aligned} 0 \leq p_z(k-1) - p_{z'}(k-1) &= p_z(k) - v_k - \big(p_{z'}(k) + u_k\big) \leq D - p_{z'}(k) - (v_k + u_k) \\ &= D - p_{z'}(k) - d_k \leq \tfrac{1}{2}p_{z'}(k) + \tfrac{1}{2}\big(D - p_{z'}(k) - d_k\big) = \tfrac{1}{2}(D - d_k) \leq \tfrac{1}{2}\delta D \ . \end{aligned}$$

Case III (see Figure 6 (b)): It remains to consider the case where both p_z and $p_{z'}$ are *backward* greedy patterns. We assume that $p_z(k-1) - p_{z'}(k-1) > \frac{1}{2}\delta D$; otherwise we are done. Then,

$$p_z(k) - p_{z'}(k) = d_k + p_z(k-1) - p_{z'}(k-1) > \big(1 - \tfrac{1}{2}\delta\big)D \ .$$

Since $p_z(m), p_{z'}(m) \in \big[\frac{1}{4}\delta D, \big(1 - \frac{1}{4}\delta\big)D\big]$ and thus $p_z(m) - p_{z'}(m) \leq \big(1 - \frac{1}{2}\delta\big)D$, there is a q with $k < q \leq m$ such that $z_q = -u_q$ and $z'_q = v_q$. Choosing the smallest such q we get

$$p_z(q-1) - p_{z'}(q-1) \geq p_z(k) - p_{z'}(k) > \big(1 - \tfrac{1}{2}\delta\big)D \ .$$

As both p_z and $p_{z'}$ are backward greedy patterns, it must hold that $p_z(q) \leq p_{z'}(q)$ since otherwise the choices $z_q = -u_q$ and $z'_q = +v_q$ contradict the backward greedy paradigm. We can thus conclude that

$$\begin{aligned} 0 \leq p_{z'}(q) - p_z(q) &= p_{z'}(q-1) + v_q - \big(p_z(q-1) - u_q\big) \\ &= d_q - \big(p_z(q-1) - p_{z'}(q-1)\big) < D - \big(1 - \tfrac{1}{2}\delta\big)D = \tfrac{1}{2}\delta D \ . \end{aligned}$$

This concludes the proof. □

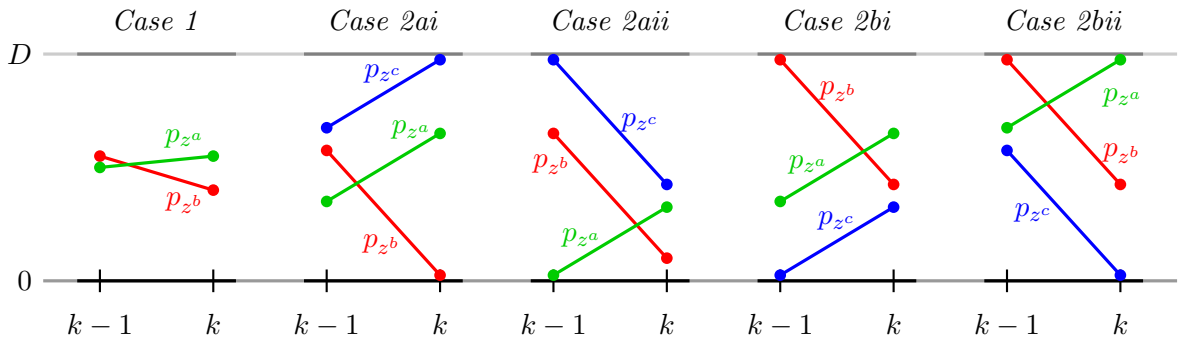

Figure 7: Illustration of different cases discussed in the proof of Lemma 4.

The next lemma identifies general situations where $\frac{1}{2}\delta D$-close patterns are guaranteed to exist.

Lemma 4. *Consider three greedy patterns $p_{z^a}, p_{z^b}, p_{z^c}$, all three living on sub-strips of $[0, D]$. If the sorting of the patterns by their end points is not a cyclic permutation of the sorting by their start points, then (at least) two of the three patterns are $\frac{1}{2}\delta D$-close.*

Proof. Since the cyclic permutations form a subgroup of the symmetric group S_3, there exists an index $k \in \{1, \dots, m\}$ such that the sorting of patterns by their values $p_{z^a}(k), p_{z^b}(k), p_{z^c}(k)$ is not a cyclic permutation of the sorting by their values $p_{z^a}(k-1), p_{z^b}(k-1), p_{z^c}(k-1)$. This means that exactly two of the three patterns, say p_{z^a} and p_{z^b}, must have exchanged their positions in the sorting when going from $k-1$ to k while the third pattern p_{z^c} has kept its position in the sorting (here we use the fact that pairwise exchanges are the only non-cyclic permutations in the symmetric group S_3).

Since p_{z^a} and p_{z^b} have exchanged their position, z_k^a and z_k^b cannot be equal, that is, without loss of generality $z_k^a = v_k$ and $z_k^b = -u_k$. Thus, $p_{z^a}(k-1) \le p_{z^b}(k-1)$ and $p_{z^a}(k) \ge p_{z^b}(k)$. An illustration of the following case distinction is given in Figure 7.

Case 1: If $d_k = u_k + v_k \le \delta D$, we show that p_{z^a} and p_{z^b} are $\frac{1}{2}\delta D$-close:

$$\begin{aligned}|p_{z^a}(k) - p_{z^b}(k)| + |p_{z^b}(k-1) - p_{z^a}(k-1)| \\ = \big(p_{z^a}(k) - p_{z^a}(k-1)\big) - \big(p_{z^b}(k) - p_{z^b}(k-1)\big) = v_k + u_k = d_k \le \delta D \ .\end{aligned}$$

This inequality implies that at position $k-1$ or k the two patterns are at most $\frac{1}{2}\delta D$ apart.

Case 2: It remains to consider the case where $d_k \ge (1-\delta)D$. With the help of Lemma 3 we argue that p_{z^c} is $\frac{1}{2}\delta D$-close to at least one of the other two patterns in this case. First notice that the unchanged position of pattern p_{z^c} in the sorting cannot be the second position between the other two patterns as $z_k^c \in \{v_k, -u_k\}$ and thus its position relative to one of the other patterns does not change when going from $k-1$ to k. Therefore, $p_{z^c}(k-1)$ and $p_{z^c}(k)$ are either both maximal or both minimal.

Case 2a: We first consider the case that $p_{z^c}(k-1)$ and $p_{z^c}(k)$ are both maximal. If $z_k^c = v_k$ (*Case 2ai* in Figure 7), then p_{z^b} and p_{z^c} are $\frac{1}{2}\delta D$-close by Lemma 3. If, on the other hand, $z_k^c = -u_k$ (*Case 2aii* in Figure 7), then p_{z^a} and p_{z^c} are $\frac{1}{2}\delta D$-close by Lemma 3.

Case 2b: The same arguments can be applied to the symmetric case where $p_{z^c}(k-1)$ and $p_{z^c}(k)$ are both minimal; see Cases 2bi and 2bii in Figure 7. □

The next lemma is the key to proving the existence of unsplittable solutions with improved additive performance (Theorem 1). From now on we set $\delta := \frac{2}{7}$, which can be seen to be optimal for our kind of analysis.

Lemma 5. *For an instance with $d_i \in \left[0, \frac{2}{7}D\right] \cup \left[\frac{5}{7}D, D\right]$ for every demand i, there is a pattern with additive performance at most $\frac{19}{14}D$.*

Proof. We start by considering a forward greedy pattern p_{z^a} with start point $x^a := \frac{5}{14}D$, living on a sub-strip of $[0, D]$. If its end point y^a is at least $\frac{2}{7}D$, then its additive performance is at most $\frac{19}{14}D$ by Observation 2 (iii) and we are done. We therefore assume that $y^a \in \left[0, \frac{2}{7}D\right]$.

Our second pattern p_{z^b} is a backward greedy pattern that ends at $y^b := \frac{3}{7}D$. Similar to above, we can assume that its start point $x^b \in \left[0, \frac{3}{14}D\right]$. If the two greedy patterns introduced so far are $\frac{1}{7}D$-close, then, by Lemma 2, there is a pattern living on a sub-strip of $\left[-\frac{1}{14}D, \frac{15}{14}D\right]$ with start point x and end point y such that $x + y = x^a + y^b = \frac{11}{14}D$. By Observation 2 (i), the additive performance of this pattern is at most

$$\max\left\{2 \cdot \tfrac{15}{14}D - \tfrac{11}{14}D, \tfrac{11}{14}D + 2 \cdot \tfrac{1}{14}D\right\} = \tfrac{19}{14}D$$

and we are done. We therefore assume in the following that p_{z^a} and p_{z^b} are not $\frac{1}{7}D$-close.

We now consider another forward greedy pattern p_{z^c} with start point $x^c := \frac{11}{14}D$. If its end point y^c is at most $\frac{4}{7}D$, then by Observation 2 (iii) its additive performance is at most $\frac{19}{14}D$ and we are done. We can therefore assume that $y^c \in \left[\frac{4}{7}D, D\right]$.

Notice that the sorting of the three patterns by their end points (i. e., a, b, c) is not a cyclic permutation of the sorting by their start points (i. e., b, a, c). Therefore, by Lemma 4 with $\delta := \frac{2}{7}$, at least two of the three patterns are $\frac{1}{7}D$-close. By our assumption above, p_{z^a} and p_{z^b} are not $\frac{1}{7}D$-close. It therefore remains to consider two cases.

If p_{z^c} and p_{z^a} are $\frac{1}{7}D$-close, then, by Lemma 2, there is a pattern living on a sub-strip of $\left[-\frac{1}{14}D, \frac{15}{14}D\right]$ with start point x and end point y such that $x + y = x^c + y^a \in \left[\frac{11}{14}D, \frac{15}{14}D\right]$. By Observation 2 (i), the additive performance of this pattern is at most

$$\max\left\{2 \cdot \tfrac{15}{14}D - \tfrac{11}{14}D, \tfrac{15}{14}D + 2 \cdot \tfrac{1}{14}D\right\} = \tfrac{19}{14}D$$

and we are done. Finally, if p_{z^c} and p_{z^b} are $\frac{1}{7}D$-close, then, by Lemma 2, there is a pattern living on a sub-strip of $\left[-\frac{1}{14}D, \frac{15}{14}D\right]$ with start point x and end point y such that $x + y = x^c + y^b = \frac{17}{14}D$. By Observation 2 (i), the additive performance of this pattern is at most

$$\max\left\{2 \cdot \tfrac{15}{14}D - \tfrac{17}{14}D, \tfrac{17}{14}D + 2 \cdot \tfrac{1}{14}D\right\} = \tfrac{19}{14}D \ .$$

This concludes the proof. □

We can finally prove our main result.

Proof of Theorem 1. If the considered instance has a demand i of value $d_i \in \left[\frac{2}{7}D, \frac{5}{7}D\right]$, the result follows by Lemma 1 with $\delta := \frac{2}{7}D$. Otherwise, the result follows from Lemma 5. □

Notice that the algorithm that is implicitly given in the proof of Theorem 1 can be implemented to run in linear time.

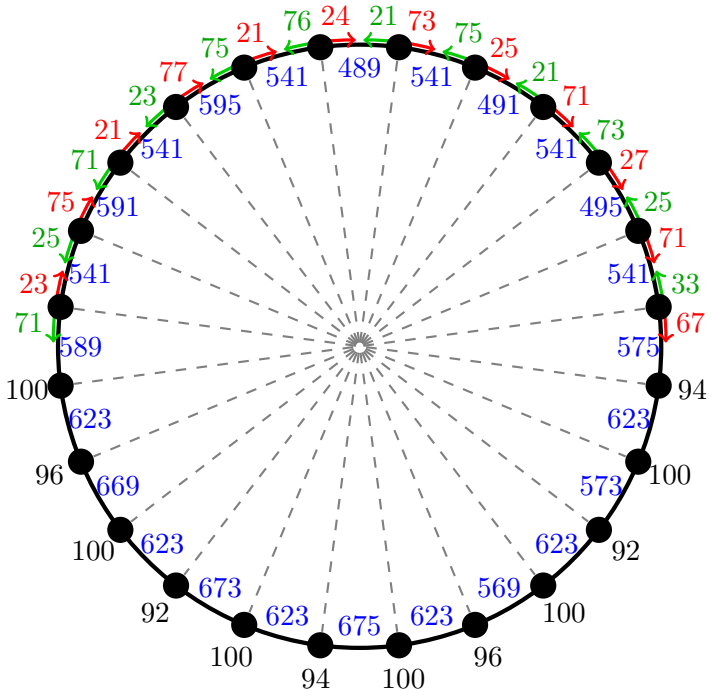

Figure 8: A slightly simplified version of the Schrijver et al. instance with $D = 100$. The given split routing cannot be turned into an unsplittable routing without increasing the load on some edge by at least 101.

5 Deriving an improved lower bound

In this section we present a counterexample to Schrijver et al.'s conjecture. While this counterexample turns out to be relatively simple and easy to verify, we would like to emphasize that finding it was a cumbersome undertaking for both the author and his computer. Schrijver et al. write in their paper [12] "We have never managed to produce a random example with $L > L^* + D \ldots$", and we can certainly confirm this experience. Among hundreds of millions of similar instances of the same size, all featuring a very special, promising structure, the one presented below turned out to be the only counterexample.

In the following we first give some more details on how we tracked down this shy beast. In fact, our primary objective was to find a split routing that cannot be turned into an unsplittable routing without increasing the load on some edge by $(1 + \varepsilon)D$ for $\varepsilon > 0$ as large as possible. Only later we realized that our insights can also be used to disprove Schrijver et al.'s conjecture.

Schrijver, Seymour, and Winkler [12] exhibit an instance of the Ring Loading Problem on a ring of $n = 26$ nodes and $m = 13$ nonzero demands together with a split routing that cannot be turned into an unsplittable routing without increasing the load on some edge by at least $\frac{101}{100}D$. An illustration of a slightly simplified version of this instance with $n = 24$ nodes and $m = 12$ demands can be found in Figure 8. The maximum demand value is $D = 100$, and it is a nontrivial task to verify that in any unsplittable routing there is an edge whose load exceeds the load of the split routing given in Figure 8 by at least 101. With the help of a computer and brute force, however, one can either check all 2^{12} unsplittable routings or, somewhat less brute, use dynamic programming

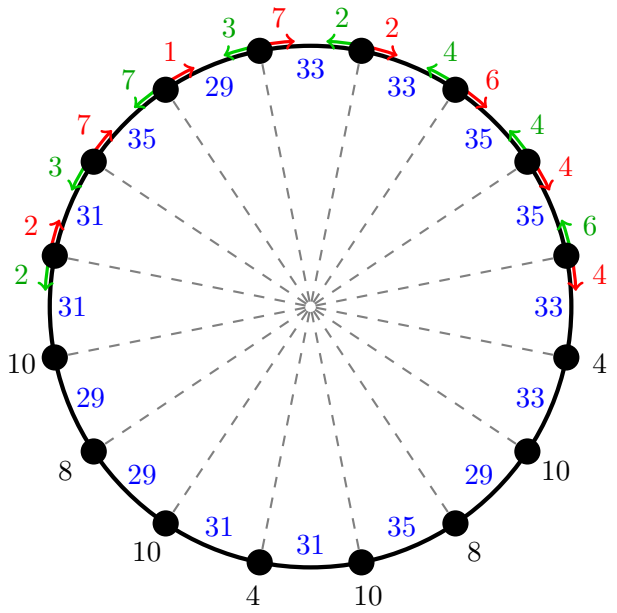

Figure 9: A split routing for 8 demands between opposite nodes together with the resulting edge loads. Each of the $256 = 2^8$ unsplittable routings increases the load on some edge by at least $11 = D + 1$.

to find an unsplittable routing with minimum possible increase of any edge load which can be done in time polynomial in n and D.[1]

In order to increase this lower bound of $\frac{101}{100}$, it is crucial to gain some understanding of the idea behind Schrijver et al.'s example: If there was an unsplittable routing z that increases the load on every edge by at most $D = 100$, it must have a corresponding pattern p_z that lives on a sub-strip of $[0, 100]$, has some start point x and 'symmetric' end point $100 - x$. Schrijver et al.'s instance, however, is constructed such that all values u_i and v_i are positive integers, $u_i + v_i$ is even, and the difference of the start and the end point of any pattern is odd, as $\sum_{i=1}^{12} u_i$ is odd. Since for integral x the difference of x and $100 - x$ is always even, the pattern would thus have to use half-integral points only and, as a consequence, lives on a sub-strip of $[0.5, 99.5]$ whose width is at most 99. This restriction of the width of suitable patterns together with the symmetry requirement on start and end points is the key property of Schrijver et al.'s instance. In this context it is important to notice that every other demand value is equal to 100 which essentially forces patterns to grow wide.

Notice that the example on $n = 16$ nodes depicted in Figure 2, with $m = 8$ demands and $D = 10$ features the same three properties as Schrijver et al.'s example: (i) all values u_i and v_i are positive integers with $u_i + v_i$ even; (ii) start and end point of any pattern differ by an odd number as $\sum_{i=1}^{8} u_i$ is odd; (iii) every other demand has value $D = 10$. For $m = 8$ and $D = 10$, there exist more than one billion such sequences $(u_1, v_1), \ldots, (u_8, v_8)$. Taking certain symmetries into account, this number can be reduced to several hundreds of millions. By brute force enumeration, we found exactly two split routings that cannot be turned into unsplittable routings without increasing the load on some edge by at least $11 = D + 1$: the one depicted in Figure 2 and the one depicted in Figure 9.

[1] The idea is to enumerate over all integral patterns living on a sub-strip of $[0, D]$ using dynamic programming.

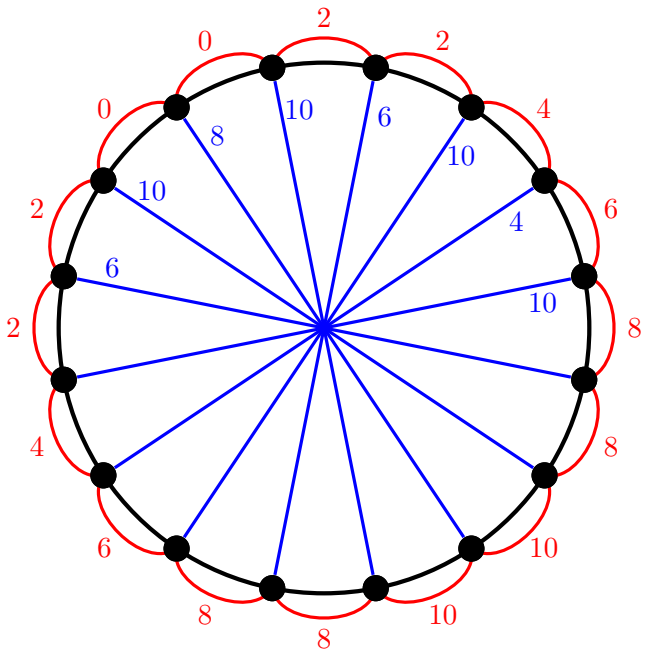

Figure 10: Extension of the instance depicted in Figure 2. By routing the additional 'short' demands unsplittably along their corresponding edges and the remaining demands between opposite nodes as suggested in Figure 2 we get an optimum split routing where all edges have load 37.

Observation 4. *There exist instances with a given split routing such that in any unsplittable solution the load on some edge is increased by at least* $\frac{11}{10}D$ *(cf. Figures 2 and 9).*

This observation does, however, not yet imply the existence of a counterexample to Schrijver et al.'s conjecture. Firstly, the split routings in Figures 2 and 9 are not optimum (notice that an optimum split routing for pairwise crossing demands splits every demand into two equal pieces). And, secondly, it is not true that any unsplittable routing increases the *maximum* edge load by more than D.

For the split routings in Figures 2 and 9 the first problem can be tackled by introducing for every edge an additional demand between its two end nodes. The demand values are chosen such that routing these demands unsplittably along their corresponding edges equalizes all edge loads; see Figure 10 for the necessary modification of the instance depicted in Figure 2. We argue that the split routing with uniform edge loads minimizes the maximum edge load. This follows from the fact that all demands are routed along shortest paths only such that the average edge load is minimum. Since the minimum and the maximum edge loads in Figure 2 (and also in Figure 9) differ by no more than D, the additional demand values can be chosen to be at most D. (Notice that this is not true for the split routing depicted in Figure 8.)

This careful modification of the considered instances almost solves the second problem as well. As a consequence of Observation 4, any unsplittable routing increases the maximum edge load by at least $\frac{11}{10}D$ as long as it sticks to routing the additional demands along their corresponding edges. The latter restriction, however, turns out to be crucial. The instance depicted in Figure 10 (based on Figure 2) has an unsplittable routing whose maximum edge load exceeds the maximum edge load of the optimum split routing by only $D = 10$.

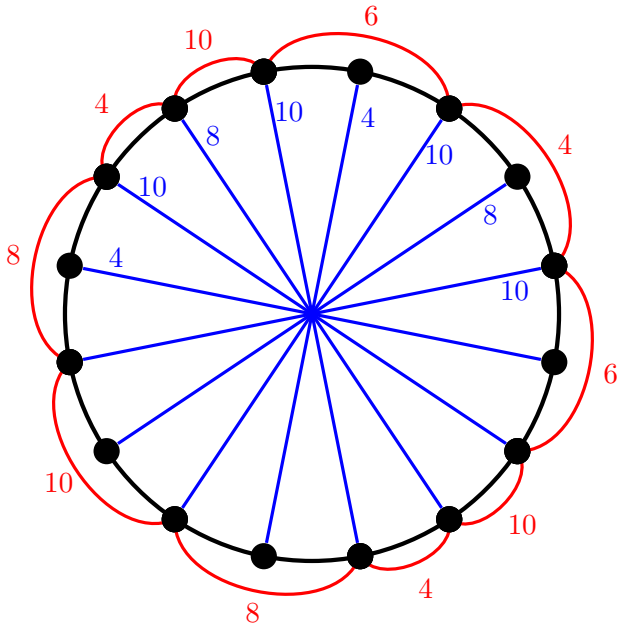

Figure 11: An instance of the ring loading problem on $n = 16$ nodes with 18 demands, 8 of which belong to opposite nodes, 4 to neighboring nodes, and 6 to nodes of distance 2. The maximum demand value is $D = 10$ and it can be checked that an optimum split routing has load $L^* = 39$ while an optimum unsplittable routing has load $L = L^* + D + 1 = 50$.

On the other hand, the extension of the instance depicted in Figure 9 does lead to a counterexample, a slightly simplified version of which is depicted in Figure 11. Similar to above, it is easy to see that combining the split routing in Figure 9 with the straightforward 'short' routing of the additional demands yields an optimum split routing where all edge loads are equal to $L^* = 39$. Finally, by enumerating all $2^{18} = 262.144$ unsplittable routings, one can verify that the optimal unsplittable load is $L = 50$ such that $L = L^* + D + 1$, thus disproving Schrijver et al.'s conjecture that $L \leq L^* + D$.

6 Conclusion

We have shown that the worst-case additive gap between optimum split and unsplittable routings for the Ring Loading Problem is at most $+\frac{19}{14}D$ and at least $+\frac{11}{10}D$. That is, $L \leq L^* + \frac{19}{14}D$ for all instances, and there exists an instance with $L \geq L^* + \frac{11}{10}D$. Our extensive yet unsuccessful search for instances yielding a larger lower bound gives us serious doubts as to whether there exist any.

Conjecture 1. *$L \leq L^* + \frac{11}{10}D$ for all instances of the* Ring Loading Problem.

We have serious doubts as to whether the algorithmic techniques and analytic tools discussed in this paper are powerful enough to close the remaining gap. But we hope that, 16 years after the publication of Schrijver et al.'s landmark paper [12], the presented progress will stimulate further research and new ideas on this fine and challenging problem.

In view of Shepherd's remark (cf. [14]) on the 'tempting similarity' of Schrijver et al.'s conjecture and results on single-source unsplittable flows [3] and, in particular, Goemans' related conjecture [2], it is natural to ask for possible consequences of the insights presented in this paper. We do not have a meaningful answer to this question, nor do we currently see any reason to doubt the validity of Goemans' conjecture.

Acknowledgements. The author would like to thank Bruce Shepherd for giving the impetus to this work as well as for numerous inspiring conversations and valuable comments.

References

[1] S. Cosares and I. Saniee. An optimization problem related to balancing loads on SONET rings. *Telecommunications Systems*, 3:165–181, 1994.

[2] G. Costain, S. Kennedy, and C. Meagher, editors. *Bellairs Combinatorial Optimization Open Problems*. 2007. www.math.mcgill.ca/ bshepherd/Bellairs/bellairs2007.pdf.

[3] Y. Dinitz, N. Garg, and M. X. Goemans. On the single source unsplittable flow problem. *Combinatorica*, 19:17–41, 1999.

[4] A. Frank. Edge-disjoint paths in planar graphs. *Journal of Combinatorial Theory, Series B*, 38:164–178, 1985.

[5] A. Frank, T. Nishizeki, N. Saito, H. Suzuki, and É. Tardos. Algorithms for routing around a rectangle. *Discrete Applied Mathematics*, 40:363–378, 1992.

[6] M. Friedrich. *Benefits of Sexual Reproduction in Evolutionary Computation*. PhD thesis, TU Berlin, 2013.

[7] T. Jansen and I. Wegener. The analysis of evolutionary algorithms—a proof that crossover really can help. *Algorithmica*, 34:47–66, 2002.

[8] S. Khanna. A polynomial-time approximation scheme for the SONET ring loading problem. *Bell Labs Technical Journal*, 2:36–41, 1997.

[9] M. Martens, F. Salazar, and M. Skutella. Representing single source multicommodity flows as convex combinations of unsplittable flows. In L. Arge, M. Hoffmann, and E. Welzl, editors, *Algorithms — ESA '07*, volume 4698 of *Lecture Notes in Computer Science*, pages 395–406. Springer, 2007.

[10] Y.-S. Myung, H.-G. Kim, and D.-W. Tcha. Optimal load balancing on SONET bidirectional rings. *Operations Research*, 45:148–152, 1997.

[11] H. Okamura and P. D. Seymour. Multicommodity flow in planar graphs. *Journal of Combinatorial Theory, Series B*, 31:75–81, 1981.

[12] A. Schrijver, P. Seymour, and P. Winkler. The ring loading problem. *SIAM Journal on Discrete Mathematics*, 11:1–14, 1998.

[13] A. Schrijver, P. Seymour, and P. Winkler. The ring loading problem. *SIAM Review*, 41:777–791, 1999.

[14] F. B. Shepherd. Single-sink multicommodity flow with side constraints. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*. Springer, 2009.

[15] M. Skutella. Approximating the single source unsplittable min-cost flow problem. *Mathematical Programming*, 91:493–514, 2002.